RL-TR-94-157
In-House Report
October 1994

PRELIMINARY EVALUATION OF A DIGITAL OPTICAL SYSTOLIC PROCESSOR IMPLEMENTED WITH ACOUSTO-OPTIC CELLS

James B. Rosetti, Rebecca J. Bussjager, David R. Christie,
1Lt, USAF, Stuart I. Libby, 1Lt, USAF, Joanne H. Maurice

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

Rome Laboratory
Air Force Materiel Command
Griffiss Air Force Base, New York

This report has been reviewed by the Rome Laboratory Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

RL-TR-94-157 has been reviewed and is approved for publication.

APPROVED: James W. Cusack

JAMES W. CUSACK, Chief
Photonics & Optics Division
Surveillance & Photonics Directorate

FOR THE COMMANDER: Donald W. Hanson

DONALD W. HANSON
Director of Surveillance & Photonics

If your address has changed, if you wish to be removed from the Rome Laboratory mailing list, or if the addressee is no longer employed by your organization, please notify RL/[illegible] Griffiss AFB NY [illegible]. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document requires that it be returned.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED

4. TITLE AND SUBTITLE

5. FUNDING NUMBERS

6. AUTHOR(S)

7. PERFORMING ORGANIZATION NAME (S) AND ADDRESS(ES)

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME (S) AND ADDRESS(ES)

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

[illegible] acousto-optic cells were imaged onto one another, and the output [illegible] to a linear photodiode array. By sending periodic data streams [illegible] processor and monitoring the output data stream, an evaluation [illegible] the potential feasibility of this type of architecture. It was shown [illegible] million logical instructions per second processor was feasible. However, [illegible] overhead in implementing higher order functions such as addition [illegible] multiplication which reduced the computational capability to 4.5 million additions [illegible] and 760 thousand multiplications per second.

14. SUBJECT TERMS	15. NUMBER OF PAGES
acousto-optic cell, systolic processor, digital [illegible] computer	28
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
[illegible]	UNCLASSIFIED	UNCLASSIFIED	UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

Contents

1. Introduction	1
2. Processor Description	4
3. System Setup	6
4. RF Drivers Characterization	8
5. Acoustic-Optic Cells Characterization	10
6. Photodiode Array Characterization	12
7. System Performance Results	14
8. Conclusions	17
References	19

1 Introduction

Much of the real time signal processing in radar applications has to be done in analog in order to achieve the required high speed. However, the accuracy in an analog circuit is directly related to its limited dynamic range. Some signal processing applications (radar mainlobe and sidelobe nulling) require accuracies that cannot be achieved via analog processing. Digital processing accuracies, however, are limited by the number of bits used in the calculations. The accuracy of a digital processor is not limited to that of its internal word length. Through the use of appropriate algorithms, an unlimited number of bits can be used in the calculations it performs, thus yielding unlimited precision (with a tradeoff in speed). The inherent parallelism of optics can be exploited to realize high speed parallel processors. This concept was proposed by OptiComp Corporation to the Office of Naval Research, contract N00014-87-C-0077[8].

This architecture is based on the ability to generate Boolean Sum of Minterms or Sum of Products (SOP) in a Systolic manner. The generation of the SOP can be done by imaging a Spatial Light Modulator (SLM) onto a second SLM to perform the logical AND operation and using a cylindrical lens to sum each column onto a linear detector. Applying a proper threshold to the output of the detector yields a logical OR operation. Figure 1-1 shows m words applied to two SLMS with height m and width n. Each word consists of n bits, and thus the detector is n elements long. The output at the detector C_n yields:

$$C_n = A_n^1 B_n^1 + \ldots + A_n^m B_n^m.$$

By rotating the cylindrical lens and the detector 90 degrees, and using a m element detector yields:

$$C_m = A_1^m B_1^m + \ldots + A_n^m B_n^m.$$

So, by rotating a lens, a different set of operations are performed. However, in this case, the logical ORing across the bits of the same word is not very useful, and would not be a practical configuration of this architecture. Note that for every data word loaded, an

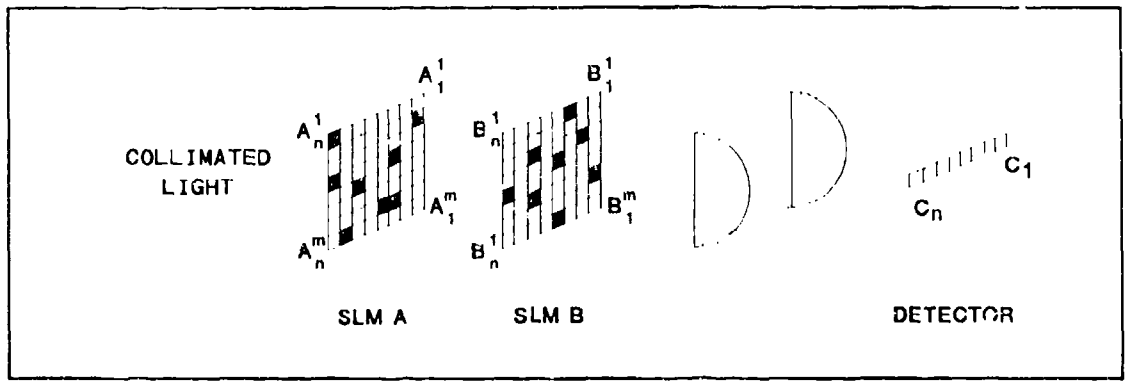

Figure 1-1: Logic Operations On Serially Loaded Data

AND followed by an OR operation is performed. If one wanted to logically OR two words together, but not do an AND operation, both words could be input to SLM A. The remainder of the SLM A would be all zeroes. SLM B would be set to all ones. To do just a logical AND of two words, one word would be input to SLM A, the other to SLM B (same row in each SLM). While this works, the true parallelism of this architecture is hardly exploited as only one operation on two binary words is carried out when the SLMs are loaded.

Another configuration which offers more parallelism is shown in Figure 1-2. Here, each pair of data words are loaded into SLM A, and instructions dictating what type

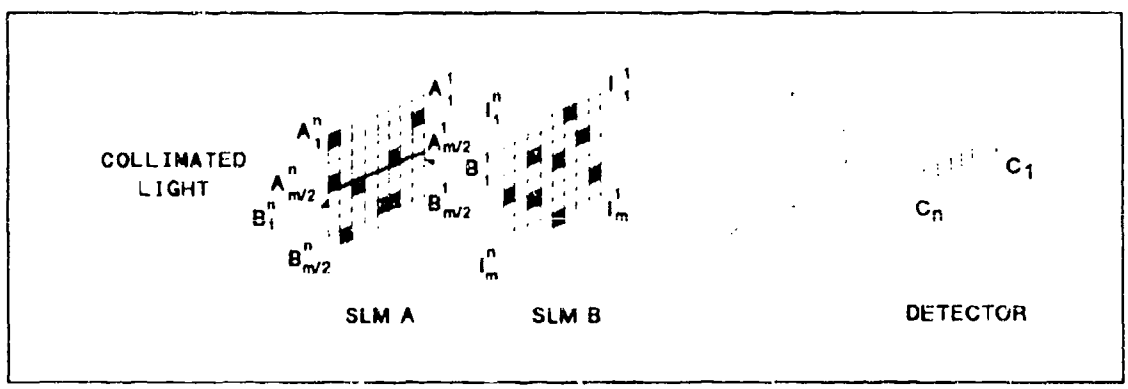

Figure 1-2: Logic Instructions Acting On Parallel Loaded Data

of operation on those words are loaded into SLM B. Thus, one could fill each SLM with data and instructions that could be done simultaneously. One drawback to this approach

is that only one logical operation is being executed--the logical OR. The second SLM only selects which bits of the two words are to be ORed. To do an AND operation, one has to make use of DeMorgan's theorem:

$$AB=\overline{\overline{A}+\overline{B}}.$$

Note this requires using the complement of the variables for input. However, this type of architecture has no capability for directly complementing bits. If an AND operation is to be performed, the data would have to be complemented before being input to the SLM. When an OR operation is required, no complementing would be done.

Systolic processing consists of a number of processors regularly interconnected. These individual processors carry out a small number of functions, typically in one clock cycle and are referred to as Processing Elements (PEs). The interconnection may be linear (see Figure 1-3), or a 2D mesh[1,2]. Data flows into a PE, and its results along with the data flow into the next PE. Since the data is also passed from PE to PE, this distinguishes a linear systolic processor from a pipelined processor[3,4]. Referring to Figure 1-3, note that although the data may enter and exit sequentially, when the processor is filled (all PEs are operating on some data simultaneously), there exists parallelism which can be quite high if there are a large number of PEs. For example, if data enters at a clock rate of 10 MHz into a linear systolic processor that has 10 PEs, the equivalent

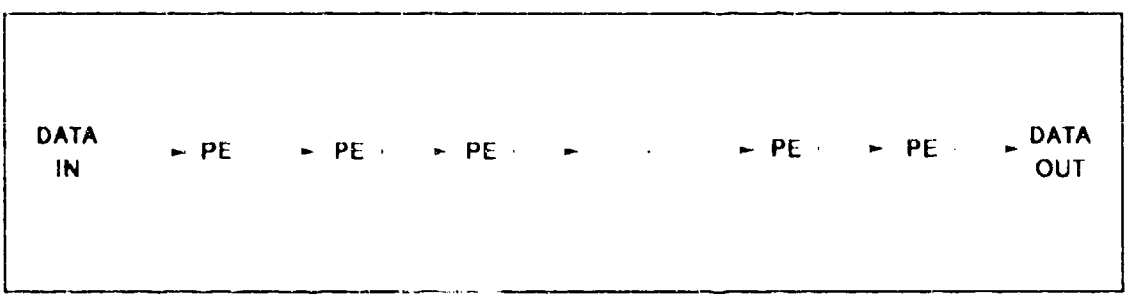

Figure 1-3: Linear Systolic Array

processing power would be that of a 100 MHz sequential processor (10 PEs working at 10 MHz each). Note two ideal conditions exist: the I/O data rate to/from the processor is low, but the overall processing bandwidth is high, and each PE consists of a low

bandwidth, simple processor. The tradeoff to this type of architecture, however, is that it is not a general purpose processor. Its applications are in signal processing such as Finite Impulse Response filters, deconvolution and matrix multiplication. Systolic processors are generally "hard wired" to a particular algorithm and thus cannot be reconfigured to perform some other task.

2 Processor Description

The Systolic Processor proposed for this evaluation consists of two multi-channel Acoustic-Optic (AO) cells for data input, and a cylindrical lens which does the logical OR function onto a linear photodiode array (see Figure 2-1). Data is input to the SLMs from

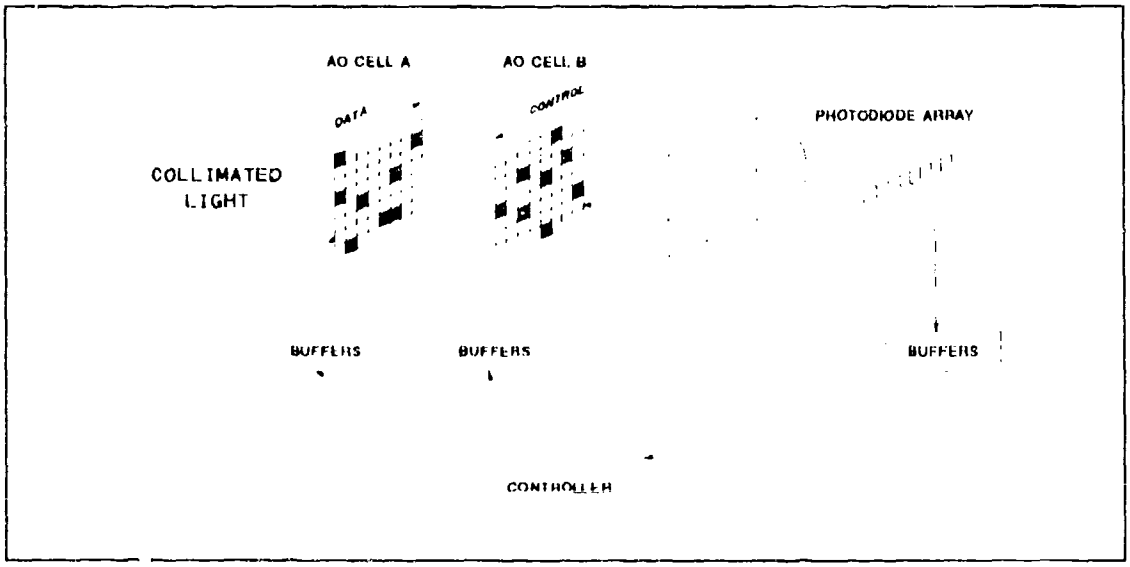

Figure 2-1: Acoustic-Optic Digital Systolic Processor

a central controller. The output data from the photodiode array is thresholded, buffered and sent back to the controller. The controller could be a general purpose electronic computer for reprogrammability, a finite state machine for speed, or a hybrid of the two.

Figure 2-2 is a diagram of one channel of the AO cell. The transducer is fed a digitally modulated RF carrier (RF pulse train). The vibrating transducer creates an acoustic wave which traverses the crystal. This acoustic wave changes the index of refraction of the material creating a phase grating which diffracts the incident light. Thus, this RF pulse train is transformed into a moving train of optical, digital data.

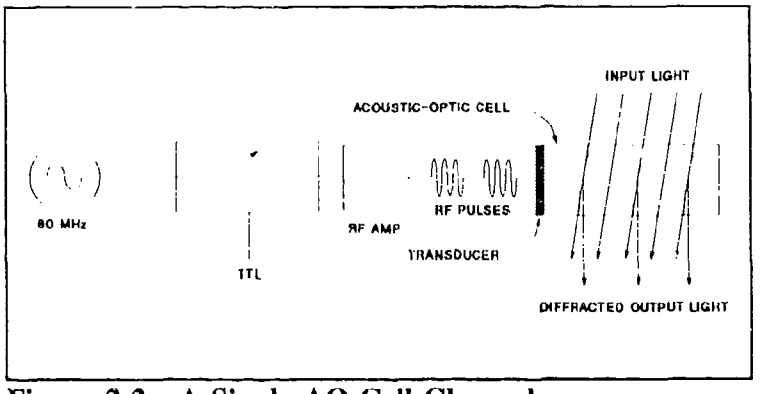

Figure 2-2: A Single AO Cell Channel

Referring back to Figure 2-1, AO cell A's diffracted optical data is imaged onto AO cell B. Thus, if a logical 1 is present on AO cell A, light will be incident on AO cell B. If, at that point in time AO cell B has a logical 1 also, the light will be diffracted and imaged onto the linear photodiode array. Data is placed in AO cell A (word parallel) and instructions in AO cell B, and they will propagate toward each other at a rate that is dependent on the type and cut of the crystal. Figure 2-3 is a block diagram of a top down view of the AO cells A and B, the cylindrical lens and the photodiode array. Data (D0 ... D4) is input to SLM A, and instructions (I0 ... I4) are entered into SLM B. At first, D0 and I0 interact. At the next time interval, D0 interacts with I1 and D1 interacts with I0. Thus, each word of the instructions serially interacts with each data word which is a form of systolic processing.

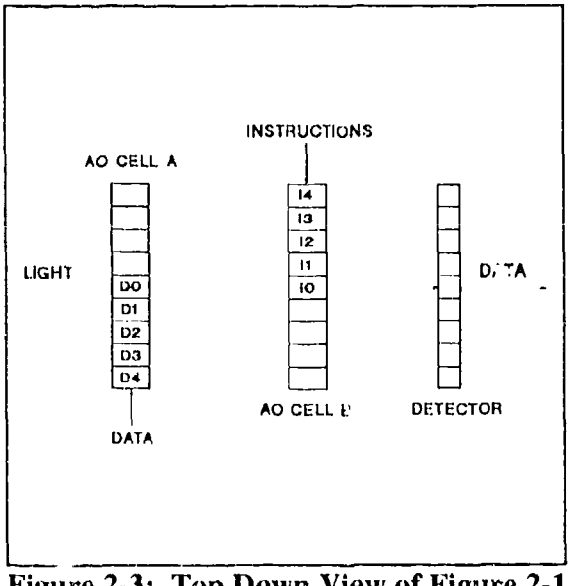

Figure 2-3: Top Down View of Figure 2-1

Recall from the Introduction, this processing architecture cannot perform negation on its own. Thus, the data is entered in dual rail format. Each bit and its complement is entered into the first AO cell so that the instructions have access to the data words and their complements simultaneously.

3 System Setup

To perform the evaluation, a subset of the processor was fabricated (see Figure 3-1). The controller and input/output buffers were omitted. Periodic waveforms from a pulse generator were used as data inputs to the SLMs to test for alignment of the system. The photodiode array was placed on a printed circuit board which was mounted on an XYZ translation stage which, in turn, was mounted to a rotation stage. RG188 coaxial cable connected each diode of the array to a patch panel of BNC connectors. This patch panel allowed probing of each photodiode, and isolated the connections from the printed circuit board so the alignment would not be disturbed. A digital scope with an internal 50 ohms resistance was connected to the patch panel and the scope's internal resistance

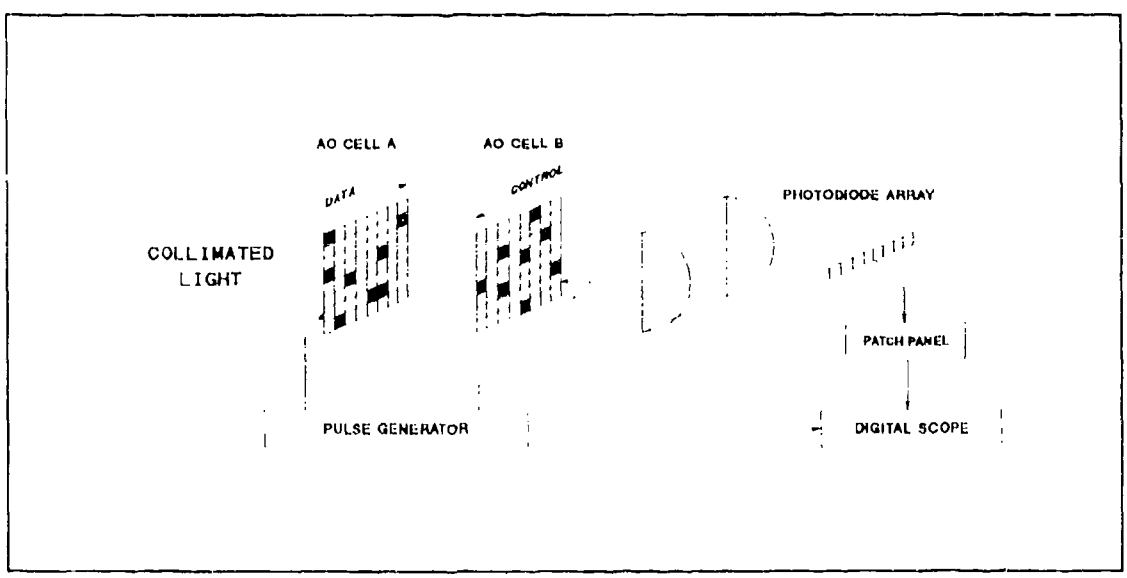

Figure 3-1: Setup For The Processor Evaluation

converted the photodiode's current into a useable voltage waveform and eliminated the need for a high speed transimpedance amplifier circuit. Load resistances of 50 ohms integrated with BNC connectors were placed on all unused diode outputs to keep the crosstalk among the photodiodes to a minimum.

A diagram of the optical layout is shown in Figure 3-2. A 10 mW HeNe laser was used as the source in the setup. The laser beam passed through collimating and expanding optics and imaged a 16 channel Acoustic-Optic cell onto another and then focused down onto the photodiode array. The first diffracted order of the AO cell was imaged onto the second AO cell. All diffracted orders from each cell, except the first needed to be filtered out (Aperture Stops). The first attempt at imaging the system involved expanding the laser to fill the entire aperture of the AO cell, but the intensity of the light incident on the photodiode array was too low to provide an adequate signal. Then, a cylindrical lens was used to focus a vertical line through the cell's aperture so each channel had a small intense beam. Since we only had a 32 element photodiode

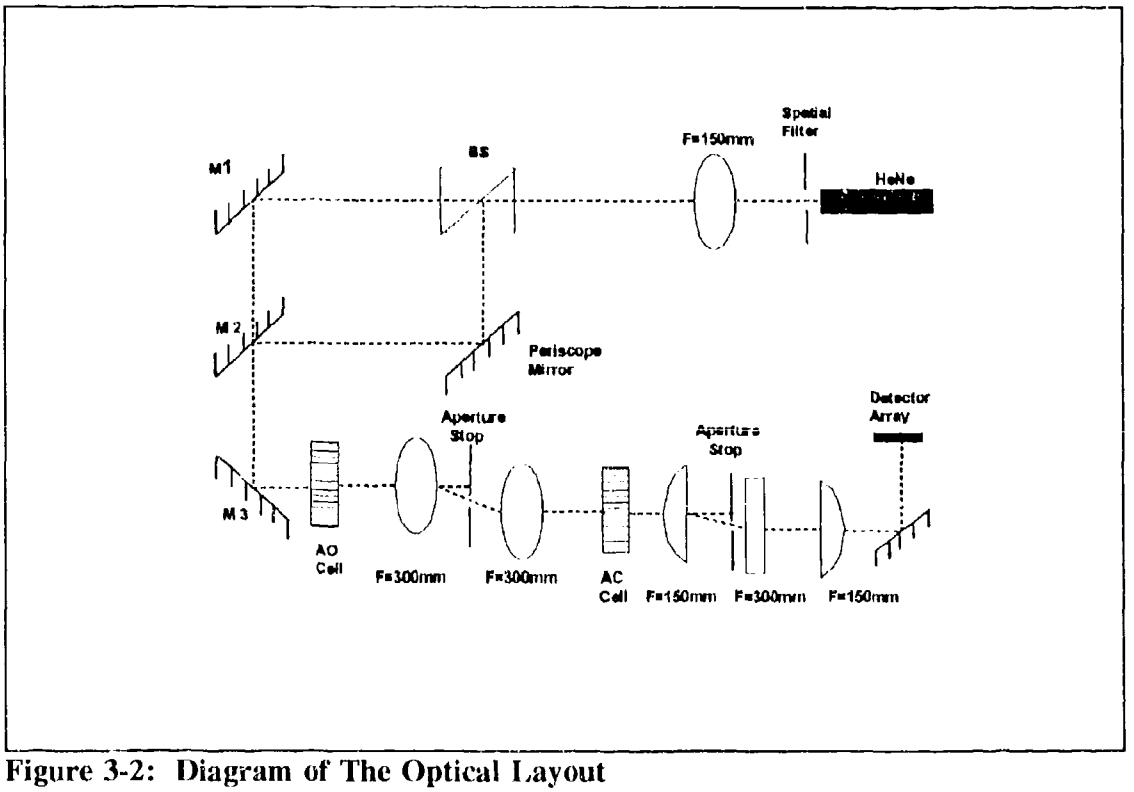

Figure 3-2: Diagram of The Optical Layout

array, there would be at most 32 bits which would not fill the entire aperture. This focused beam would yield more optical power incident on each pixel of the photodiode array. Focusing a line on the cell and using a telescope to image this line onto the second cell, however, decreased the resolution between the zero and first diffracted order (smeared them together) to the point where this setup was not useable.

Instead, the laser was split into two beams illuminating the upper and lower channels of the AO cells. This would not illuminate all channels of the AO cells. A fully illuminated system would require multiple lasers illuminating different channels of the AO cells. It was felt this would sufficiently test the capabilities of the optics to adequately focus the light onto the detector array as more distortion occurs the farther an image is from the optical axis (which passes through the center of the AO cells). Therefore, channel 16 on the first AO cell was lined up with channel 15 on the second cell since channel 16 on the second cell was not functioning within specifications. Similarly, channel 2 on the first cell (channel 1 was below specs) was lined up with channel 1 on the second cell. In order to provide light to two channels, the laser beam was collimated and split by a beamsplitting cube. The beam that was split reentered the optical train at a lower level via mirror M2.

4 RF Drivers Characterization

The RF drivers were manufactured by Newport Electro-Optics Systems. Each driver is capable of delivering 300 mW at a center frequency of 80 MHz. The output is switchable between Continuous Wave, CW, and externally modulated. The input to the external modulation is TTL compatible.

The CW output power of each driver was measured using an RF spectrum analyzer and a HP 8496B attenuator (to keep the analyzer's front end from being overloaded). Almost all outputs had to be adjusted to bring them within specifications. The output power of each channel also drifted after a day, but not enough to affect the experiment. To be safe, an RF power meter was used to monitor the output power levels.

The extinction ratio was obtained by calculating the ratio of the output power when the driver was switched from CW to off. The off state was obtained by switching the driver to external modulation with a logic 0 input to the external modulation port. The average value found was 65 dB which was more than adequate for this experiment.

The spectral purity of the drivers was checked by setting the driver to CW, and applying the output to an RF attenuator connected to an RF spectrum analyzer and measuring the amplitude of the first harmonic frequency relative to the center frequency. Poor spectral purity will result in added harmonics distorting the sinusoid (see Figure 4-1). The average value was -17.5 dBc (dB below the center frequency).

The risetime of the RF pulses was measured by modulating the driver, and applying the output to an RF attenuator connected to a digital scope. The inverted output of the pulse generator was used to trigger the scope so the edges of the modulation envelope would be observed (otherwise, the scope would trigger on the RF). The display of the scope was set to envelope mode which displays the minimum and maximum points of a waveform. After a couple of seconds, the modulation envelope of the RF pulse would appear because the unsynched 80 MHz frequency contained in the envelope would trace it out. The pulse generator's output is designed to drive a 50 ohm load, but because of the TTL load of the RF driver it is connected so there would be an additive reflection at the input of the RF driver. Therefore, the amplitude of the pulses was set to approximately 2.2V which, when reflected, became approximately 4.4V at the input of the driver. Later, when the BNC tee connectors were used to drive extra channels, the pulses were degraded because the

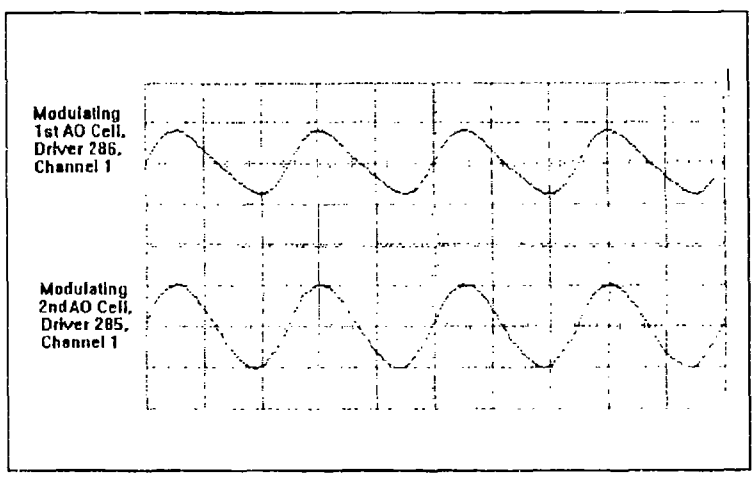

Figure 4-1: RF Waveforms From The Drivers

reflections from each load (the drivers) were not arriving back at the pulse generator at the same time. Instead, a 50 ohm load was connected to each driver input. Although this would present a 25 ohm load to the pulse generator, the reflections were minimal, and better shaped pulses from the pulse generator were obtained.

Four drivers were selected which had the best spectral purity and rise/fall time characteristics. The drivers selected were channels 6 and 7 on driver 285, and channels 4 and 7 on driver 280. Each channel's envelope was stored in the digital scope's memory, and the envelopes of the selected drivers are shown in Figures 4-2 and 4-3.

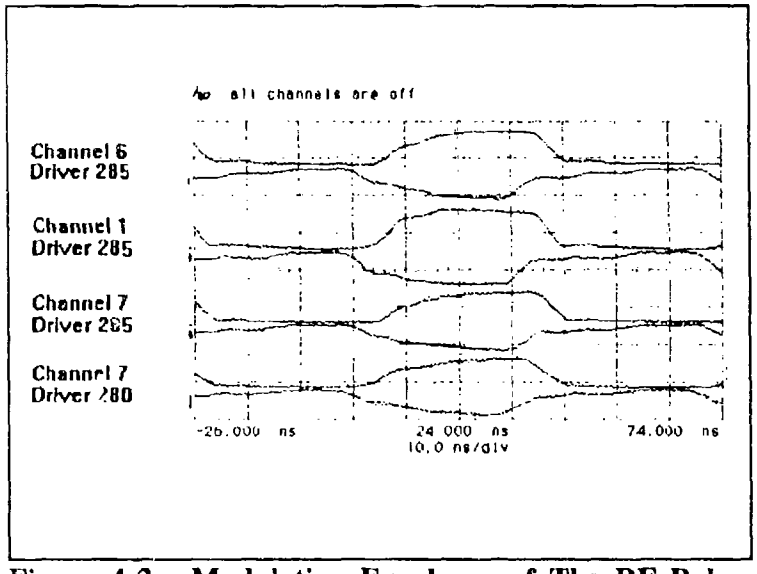

Figure 4-2: Modulation Envelopes of The RF Pulses From The Drivers

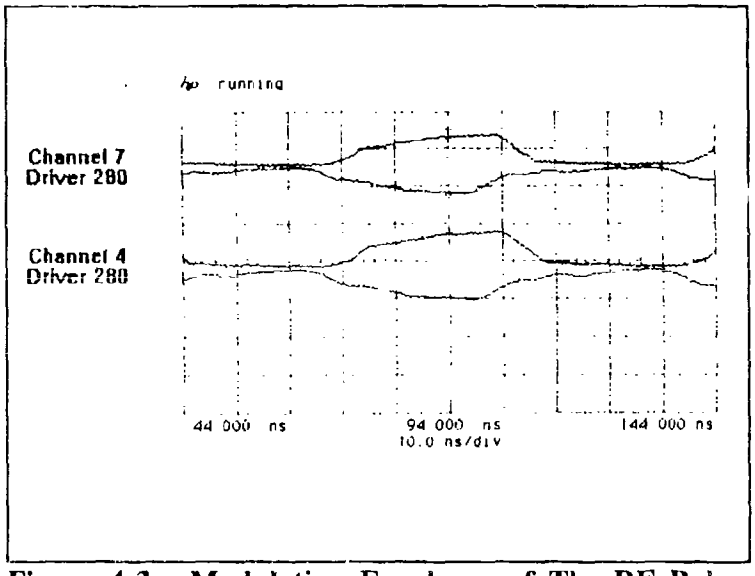

Figure 4-3: Modulation Envelopes of The RF Pulses From The Drivers

5 Acoustic-Optic Cells Characterization

The AO cells are 16 channel cells manufactured by Newport EOS. The material is Tellurium Oxide. The center frequency for the cells is 80 MHz and the optimum

wavelength is 680 nm. The Lithium Niobate transducers have a length and width of 16 mm and 0.9 mm respectively. Their spacing is 2.5 mm resulting in an overall height of 45 mm for the cell. The cell's width is 12 mm.

The diffraction efficiency of each channel of the AO cells was measured using the setup in Figure 5-1. The cell is rotated to the Bragg angle with respect to the incident laser which produces the maximum amount of diffraction into the first order. With no RF applied to the input of the channel being measured (no diffraction), the resulting optical power was measured. Then, RF power was applied and the resulting first order optical power was measured. The efficiency is:

$$Efficiency = \frac{First\ Order\ Optical\ Power}{Nondiffracted\ Optical\ Power}\ x\ 100\%$$

The average diffraction efficiencies for each AO cell are 71% for AO cell 0050 and 65% for AO cell 0051.

The propagation velocity of the acoustic wave travelling across the cell was measured by using the setup shown in Figure 3-1. Arbitrarily long pulses from the pulse generator were used to modulate a channel of one AO cell while the driver to the other cell was set to CW. Three consecutive detectors were observed on the digital scope. The difference in delay from adjacent photodiodes corresponds to the time it took for the acoustic wave to propagate from one photodiode to the other (see Figure 5-2). Since the photodiode width is known (150 μm), the propagation velocity is this width divided by the delay time (34.5 ns). This resulted in a propagation velocity of 4.3 km/s which compares favorably with the specified 4.26 km/s.

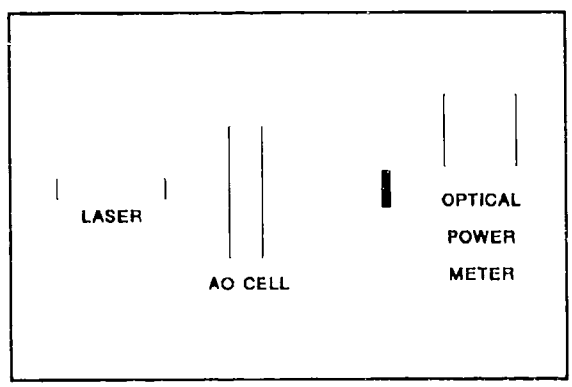

Figure 5-1: Setup For Measuring Diffraction Efficiency

6 Photodiode Array Characterization

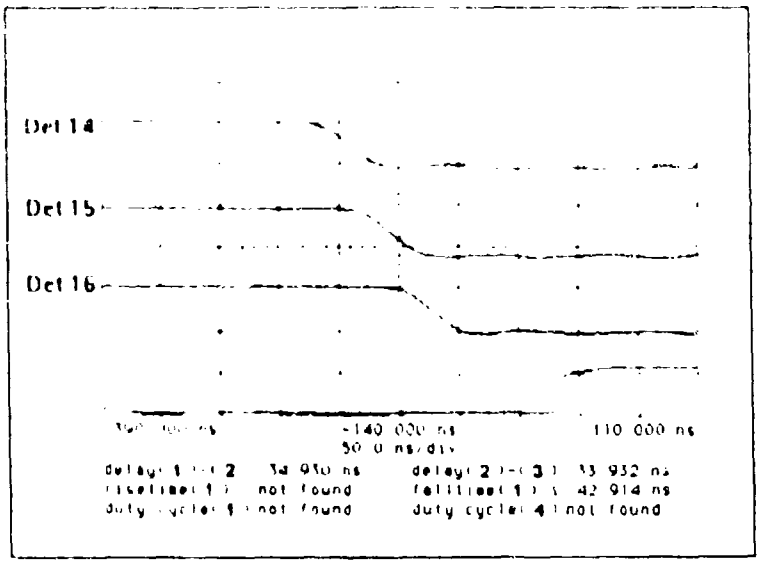

Figure 5-2: Acoustic Wave Delay

The photodiode array is a 32 element silicon avalanche photodiode array manufactured by RCA, part number C30631E. The element size is 150 μm wide by 500 μm high. Its specified rise/fall times are 2 ns. The array is lensed which results in minimal dead space between pixels[5]. Its operating voltage is 300 V.

The rise/fall times of the detectors were measured by applying optical pulses with rise/fall times of approximately 1 ns from the York laser diode to the detector array. The output of the array was applied to the 50 ohm input of the digital scope. The average rise/fall times measured were 2 ns.

The uniformity of the responsivity of the photodiode array was checked by applying a uniform intensity light source across the entire array, and measuring the resultant output current. To eliminate the Guassian intensity profile, a HeNe laser was expanded from a diameter of 0.68 mm to approximately 4.4 cm and collimated using a 65X microscope objective and a 10μm pinhole. The photodiode array was reverse-biased with 200V and all elements not being measured were terminated with 50 ohm loads. To test a particular detector, the 50 ohm load

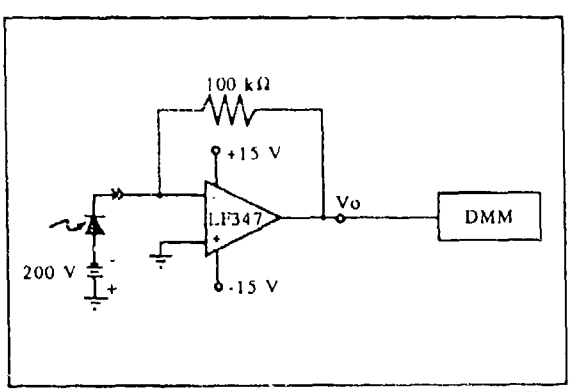

Figure 6-1: Transimpedance Amplifier Circuit

was removed from that element, a transimpedance amplifier (see Figure 6-1) with 100k transimpedance was inserted in its place to convert the photodiode's output current to a voltage, and the amplifier's output voltage was measured with a Fluke digital multimeter. This method was repeated for each of the elements.

The output signal for all 32 elements was measured three times and averaged, resulting in a graph of the output voltage, Vo, versus Element # (see Figure 6-2). Note a sinusoidal response is noticeable across the photodiode array. The light source was checked for uniformity by monitoring the output of a single photodiode, and translating the entire array horizontally across the incident laser beam. Fluctuations in the amplitude of the photodiode current were observed. Also, a laser beam analyzer showed interference patterns which were caused by reflections from the lenses in the optical system. Element #4 showed a greater response than other elements. Further study showed that this element produced current when the outside edge of element #1 was illuminated with a fiber light source. This was attributed to reflections occurring inside the package of the photodiode array. This also occurred in the other photodiode array, and was more pronounced. The affected elements were not needed for this experiment.

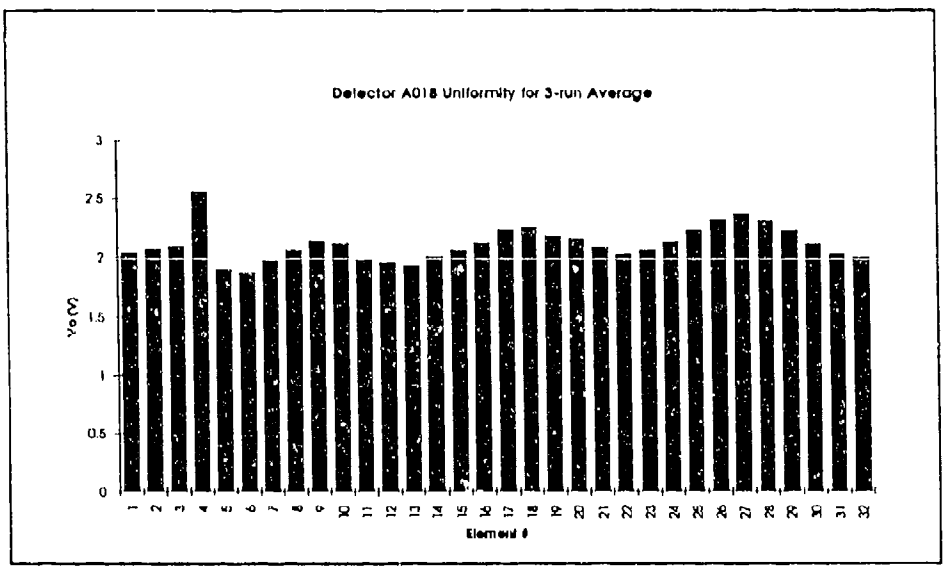

Figure 6-2: Average Responsivity of The Photodiode Array

7 System Performance Results

To evaluate the performance potential of this architecture, the data input rate to the system would be increased until the system experienced crosstalk from adjacent photodiodes. The system data input used were two repetitive waveforms. The first waveform was a square wave. Note the output is analogous to an analog convolution of the two inputs to the AO cells and the photodiode array because of the fact the two inputs propagate towards each other and across the photodiode array. Referring to Figure 7-1, the expected output when one AO cell is CW and the other has a square wave is a repetitive triangle wave. With a square wave on both AO cells, the output will be a triangle pulse waveform shown in Figure 7-2. If one of the inputs does not align properly in time with the other input and the photodiode array, a trapezoidal waveform will result. Note once the proper waveforms are obtained, this will show only vertical alignment of the system. The alignment in the horizontal direction, perpendicular to the optical axis, may be off an integer number of elements in the photodiode array. To achieve this alignment, a different input waveform is

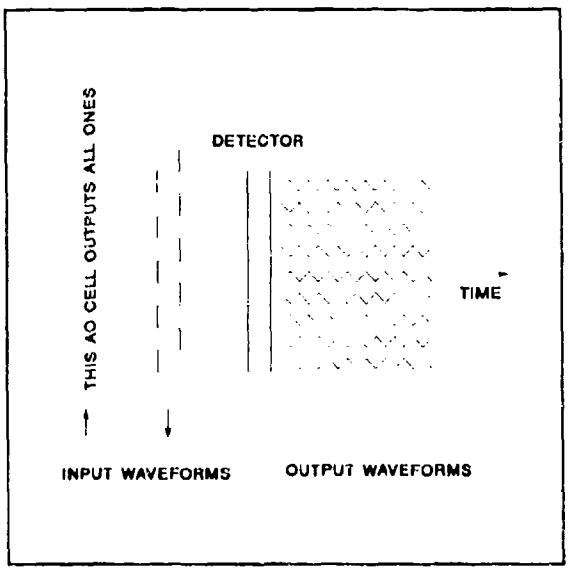

Figure 7-1: Expected Output Waveforms When One AO Cell is Input CW And The Other a Square Wave

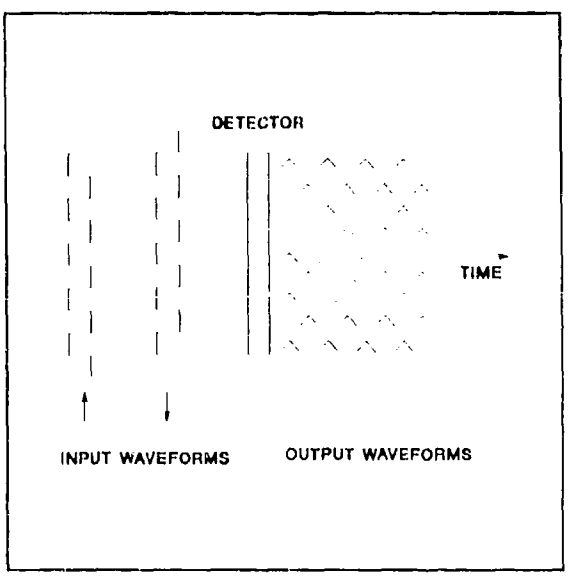

Figure 7-2: Expected Output Waveforms When Both Inputs Are Square Waves

required. The waveform is shown in Figure 7-3 with its corresponding expected output waveform. Once the previous alignment is obtained, this alignment can be obtained by either translating one of the AO cells or delaying one of the inputs to the AO cells relative to the other input (using the delay capability of the pulse generator). Since nothing is to be gained by this alignment (the crosstalk shows up in the previous alignment), it was not done in this experiment.

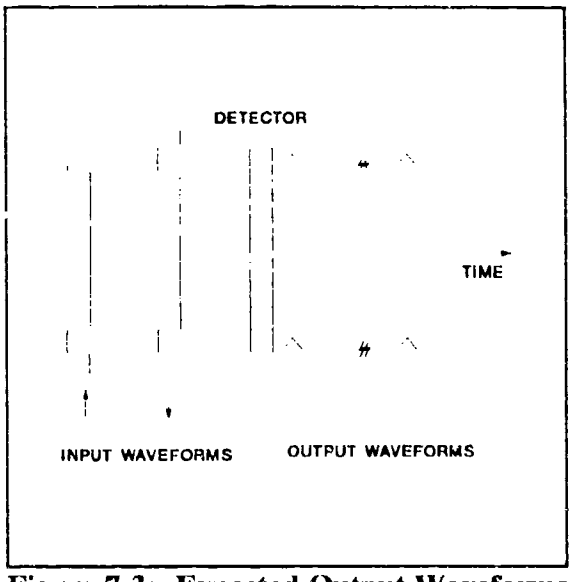

Figure 7-3: Expected Output Waveforms When Both Inputs To The AO Cells Are Repetitive Pulses

The alignment proceeded in two phases. In the first phase, only a single beam passed through the center of the optics, keeping distortion (and the number of variables) to a minimum. During this alignment, other diffraction orders were reaching the photodiode array. The patterns would change when the pulse widths of the data inputs to the AO cells were changed. These orders were due to a second order grating being established by the repetitive spacing of the phase grating in the AO cell. Figure 7-4a shows a typical channel of an AO cell. The vertical dotted lines depict the acoustic wave propagating across the crystal. When the spacing between the input RF pulses to the AO cell become small enough, light begins to diffract (see Figure 7-4b). This was verified by observing the diffracted output beam with a laser beam analyzer while varying the pulse widths of the data. The analyzer showed a diffraction order "spot" move as the pulse width was changed. The filter slit was narrowed to remove these orders.

Once satisfactory waveforms were obtained (see Figure 7-5), two beams illuminating the upper and lower channels of the AO cells were used. Note the triangle pulse points down as opposed to the waveform shown in Figure 7-1. This is because of the reverse bias of the photodiodes. The data rate was high enough to produce crosstalk in the output once the system was aligned. Figure 7-6 shows the resulting output

waveforms. Note the center of the zero has a small triangle (a "one") indicating the crosstalk. There was not enough optical power incident on the edges of the photodiode array to generate observable waveforms. Thus, only 21 of the 32 elements of the photodiode array were able to be observed. The input data rate to each channel of the AO cells was approximately 30 Mb/s.

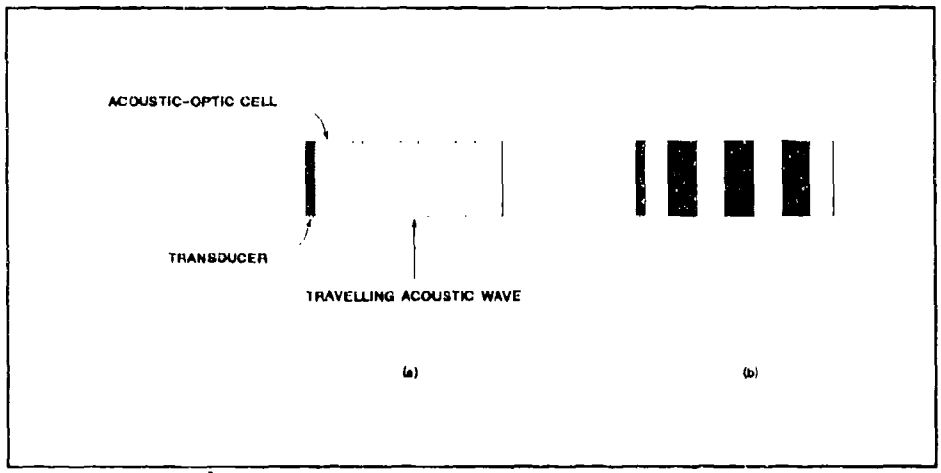

Figure 7-4: 2nd Order Diffraction Grating

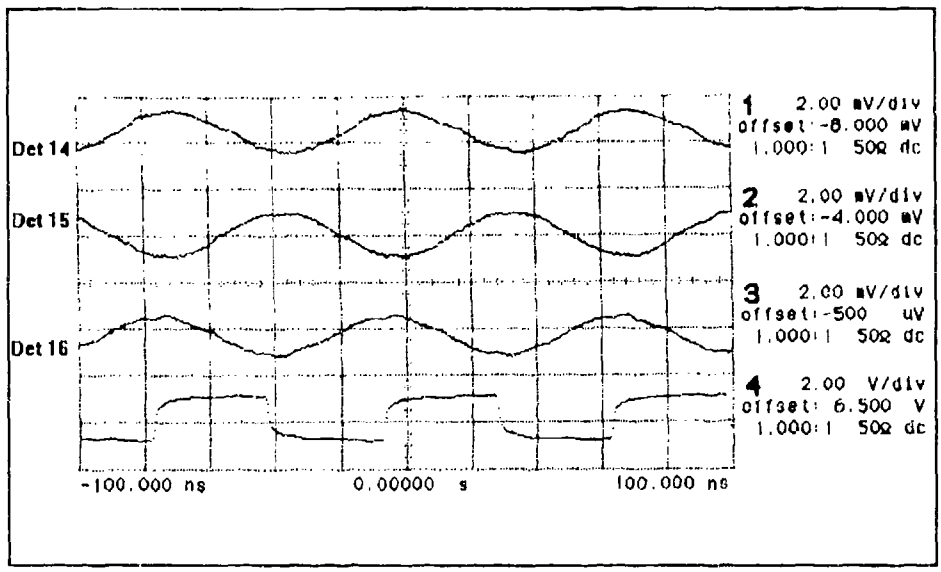

Figure 7-5: Output Waveforms With One AO Cell Input CW, The Other Modulated By a Square Wave

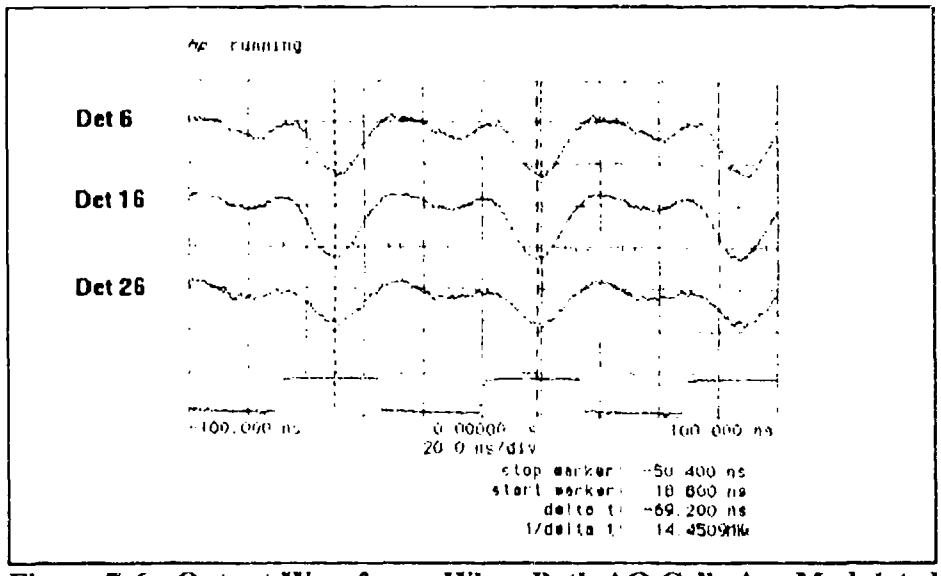

Figure 7-6: Output Waveforms When Both AO Cells Are Modulated By Square Waves

8 Conclusions

The data rate used in this experiment required bit widths on the order of 35 ns. Recall the RF pulse waveforms shown in Figures 4-2 and 4-3. It is clear we were stressing the capabilities of the RF drivers. The trailing and leading edges of the pulses are not only poor, they are very different from each other. Furthermore, the modulation envelopes of the four RF pulses are not identical. This accounts for the alignment difficulties experienced when two channels of the AO cells were illuminated. It is evident the RF drivers play an important role in the success of this type of architecture.

The instructions for this processor are low level Boolean instructions which determine what logical OR or logical AND operations are to be performed. They are not the high level add or multiply instructions which are used in electronic processors. Therefore, it would not be appropriate to refer to the computational capability of this processor as MIPS. Since 21 elements were aligned, this results in an overall computational capability of 630 Million Logical Instructions Per Second (MLIPS).

To perform an n bit by n bit addition using Full Adders (F/As) with ripple carry requires a Half Adder (H/A) plus n-1 F/As[6]. A H/A requires one EXOR gate and an AND gate. Since this processor cannot do an EXOR operation, the EXOR gate can be expanded to its equivalent form using AND and OR gates, increasing the gate count to 4. Similarly, it takes 9 logic gates to perform the function of a F/A. Therefore, it takes 4 + 9(n-1) logic gates to perform the addition. Thus, for a 16 bit addition, 139 logic operations are required. With an overall computational capability of 630 MLIPS, the processor could do 4.5 million additions per second (630 MLIPS/139 gates). The overhead for multiplication is worse. For an n bit by n bit multiplication done by forming partial products and using F/As to sum these products to produce the answer requires n^2 AND gates to generate the partial products and n(n-1) F/As[7]. This results in a total gate count of 2,416 logic gates for a 16 bit multiplications which yields 260 thousand multiplies per second which is hardly impressive. Note these numbers are assuming 100% efficiency in the utilization of the processor (it is filled at all times).

A potential bottleneck to this processor is the I/O to/from it. To make this system workable, amplifiers and thresholding logic and buffers would be required at the output. The detector has a -3 dB frequency cutoff of 175 MHz (0.35/risetime). This equates to a maximum data rate of 350 Mb/s. At data rates greater than this, the necessary buffer electronics and photodiode amplifiers/thresholding logic becomes expensive and difficult. At this data rate, to perform 100 million 16 bit multiplications per second would require AO cells with 64 channels (2 channels per bit for the dual rail and placing both the multiplier and the multiplicand on the AO cell). The processor would be required to process at a rate of 2.416×10^5 MLIPS (2,416 operations per multiply times 100 million multiplies per second). To achieve this, a photodiode array with 690 elements is required (2.416×10^5 MLIPS divided by 350 Mb/s per photodiode). This, of course, would require buffering 690 data lines at a clock rate of 350 MHz which is impractical. To lower the number of data lines from the output to a more manageable number of 64 would require the increase of the data rates by over an order of magnitude.

References

1. Kung, H. T., "Why Systolic Architectures?", Computer, pp. 37-46, January 1982.

2. Li, Guo-Jie and Wah, Benjamin W., "The Design of Optimal Systolic Arrays", IEEE Transactions On Computers, Vol. C-34, No. 1, pp. 66-77, January 1985.

3. Stone, Harold S., High-Performance Computer Architecture, Addison-Wesley Publishing Company, Reading, MA, 1987.

4. Kogge, Peter M., The Architecture of Pipelined Computers, Hemisphere Publishing Corporation, New York, NY, 1981.

5. Trakalo, Murray, Webb, Pierre Poirier, McIntyre, Robert, "Avalanche photodiode thirty-two-element linear array with minimal dead space", Applied Optics, Vol. 26, No. 17, pp. 3594-3599, September 1, 1987.

6. Fletcher, William I., An Engineering Approach To Digital Design, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1980.

7. Waser, Shlomo, "High-Speed Monolithic Multipliers for Real-Time Digital Signal Processing", Computer, pp. 19-29, October 1978.

8. OptiComp Corp., "Optical Digital Computer Architecture Development and Design Specification Interim Technical Report", OptiComp Technical Report, March 1988.

*U.S. GOVERNMENT PRINTING OFFICE 1990-610-[illegible]/00077

MISSION

OF

ROME LABORATORY

Mission. The mission of Rome Laboratory is to advance the science and technologies of command, control, communications and intelligence and to transition them into systems to meet customer needs. To achieve this, Rome Lab:

a. Conducts vigorous research, development and test programs in all applicable technologies;

b. Transitions technology to current and future systems to improve operational capability, readiness, and supportability;

c. Provides a full range of technical support to Air Force Materiel Command product centers and other Air Force organizations;

d. Promotes transfer of technology to the private sector;

e. Maintains leading edge technological expertise in the areas of surveillance, communications, command and control, intelligence, reliability science, electro-magnetic technology, photonics, signal processing, and computational science.

The thrust areas of technical competence include: Surveillance, Communications, Command and Control, Intelligence, Signal Processing, Computer Science and Technology, Electromagnetic Technology, Photonics and Reliability Sciences.